AGU abstract

Systematic comparison of automated geological feature detection methods for impact craters

Accurate, automated crater counts will be essential in extrapolating from existing Mars crater catalogs to much larger catalogs of impact craters in high-resolution orbital imagery for use in relative dating of surfaces in such imagery. Once validated, automatic methods for performing crater counts could be integrated into tools such as the Planetary Image Atlas, which is designed to be a convenient interface through which a user can search for, display, and download images and other ancillary data for planetary Missions, and the Diamond Eye image mining system.

Here we report on preliminary computational experiments in using a trainable feature detection algorithm [Burl et al. 2001] to detect craters in real and simulated Mars orbital imagery, and to derive approximate impact crater counts for geological use. In these experiments, we consider two uses of the trainable feature detector: first, directly as a crater detector, and second, as two detectors for sunlit and shadowed inner walls of craters which can then be assembled into a single crater detection based on multiple pieces of evidence. For both of these methods, we consider two data sources: one consisting of real Viking Orbiter imagery of Mars with human expert-supplied ground truth labels, and the other consisting of computer generated renderings of simplified, synthetic cratered terrain with 100% accurate ground truth labels and known, controllable crater density. Each detector reports out a numeric detection "likelihood" for every candidate crater. This likelihood must then be thresholded to produce a detection decision.

For each combination of two data sources (one natural and one synthetic) and two crater detection methods (whole-crater and parts-model), we vary image complexity and finally measure detection accuracy. Detection accuracy is measured by a Receiver Operator Characteristic (ROC) curve in which detection efficiency (the fraction of true craters detected) and purity (the fraction of detected craters which are also true craters) are plotted against one another as a control parameter is varied, namely the likelihood threshold for deciding that a feature has been detected. The results allow a comparison of alternative geological feature-detection algorithms and show their relative strengths and weaknesses, and directions for future improvement. We also plot purity as a function of likelihood threshold in order to recalibrate the detection algorithm's own estimate of its accuracy. Finally we measure the accuracy with which an imperfect detector can be used to estimate true crater counts in an image, as a function of image complexity.